

[March-2021]Free Braindump2go DP-203 65 DP-203 Dumps Free Download[Q46-Q61]

March/2021 Latest Braindump2go DP-203 Exam Dumps with PDF and VCE Free Updated Today! Following are some new DP-203 Real Exam Questions!

Question: 46 You configure monitoring for a Microsoft Azure SQL Data Warehouse implementation. The implementation uses PolyBase to load data from comma-separated value (CSV) files stored in Azure Data Lake Gen 2 using an external table. Files with an invalid schema cause errors to occur. You need to monitor for an invalid schema error. For which error should you monitor?

A. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge_Connect: Error[com.microsoft.polybase.client.KerberosSecureLogin] occurred while accessing external files.'

B. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge_Connect: Error [No FileSystem for scheme: wasbs] occurred while accessing external file.'

C. Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11": for linked server "(null)", Query aborted-- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed.

D. EXTERNAL TABLE access failed due to internal error: 'Java exception raised on call to HdfsBridge_Connect: Error [Unable to instantiate LoginClass] occurred while accessing external files.'

Answer: C

Explanation: Customer Scenario: SQL Server 2016 or SQL DW connected to Azure blob storage. The CREATE EXTERNAL TABLE DDL points to a directory (and not a specific file) and the directory contains files with different schemas. SSMS Error: Select query on the external table gives the following error: Msg 7320, Level 16, State 110, Line 14 Cannot execute the query "Remote Query" against OLE DB provider "SQLNCLI11" for linked server "(null)". Query aborted-- the maximum reject threshold (0 rows) was reached while reading from an external source: 1 rows rejected out of total 1 rows processed. Possible Reason: The reason this error happens is because each file has different schema. The PolyBase external table DDL when pointed to a directory recursively reads all the files in that directory. When a column or data type mismatch happens, this error could be seen in SSMS. Possible Solution: If the data for each table consists of one file, then use the filename in the LOCATION section prepended by the directory of the external files. If there are multiple files per table, put each set of files into different directories in Azure Blob Storage and then you can point LOCATION to the directory instead of a particular file. The latter suggestion is the best practices recommended by SQLCAT even if you have one file per table. Incorrect Answers: A: Possible Reason: Kerberos is not enabled in Hadoop Cluster. References: <https://techcommunity.microsoft.com/t5/DataCAT/PolyBase-Setup-Errors-and-Possible-Solutions/ba-p/305297>

Question: 47 You use Azure Data Lake Storage Gen2. You need to ensure that workloads can use filter predicates and column projections to filter data at the time the data is read from disk. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Reregister the Microsoft Data Lake Store resource provider.

B. Reregister the Azure Storage resource provider.

C. Create a storage policy that is scoped to a container.

D. Register the query acceleration feature.

E. Create a storage policy that is scoped to a container prefix filter.

Answer: BD

Question: 48 DRAG DROP You plan to monitor an Azure data factory by using the Monitor & Manage app. You need to identify the status and duration of activities that reference a table in a source database. Which three actions should you perform in sequence? To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
From the Data Factory monitoring app, add the Source user property to the Activity Runs table.	
From the Data Factory monitoring app, add the Source user property to the Pipeline Runs table.	
From the Data Factory authoring UI, publish the pipeline.	
From the Data Factory monitoring app, add a linked service to the Pipeline Runs table.	
From the Data Factory authoring UI, generate a user property for Source on all activities.	
From the Data Factory authoring UI, generate a user property for Source on all datasets.	

Answer:

From the Data Factory authoring UI, generate a user property for Source on all activities.

From the Data Factory monitoring app, add the Source user property to the Pipeline Runs table.

From the Data Factory authoring UI, publish the pipelines.

Step 1: From the Data Factory authoring UI, generate a user property for Source on all activities. Step 2: From the Data Factory monitoring app, add the Source user property to Activity Runs table. You can promote any pipeline activity property as a user property so that it becomes an entity that you can monitor. For example, you can promote the Source and Destination properties of the copy activity in your pipeline as user properties. You can also select Auto Generate to generate the Source and Destination user properties for a copy activity. Step 3: From the Data Factory authoring UI, publish the pipelines. Publish output data to data stores such as Azure SQL Data Warehouse for business intelligence (BI) applications to consume. References:

<https://docs.microsoft.com/en-us/azure/data-factory/monitor-visually> Question: 49 You have an enterprise data warehouse in Azure Synapse Analytics named DW1 on a server named Server1. You need to verify whether the size of the transaction log file for each distribution of DW1 is smaller than 160 GB. What should you do? A. On the master database, execute a query against the sys.dm_pdw_nodes_os_performance_counters dynamic management view. B. From Azure Monitor in the Azure portal, execute a query against the logs of DW1. C. On DW1, execute a query against the sys.database_files dynamic management view. D. Execute a query against the logs of DW1 by using the Get-AzOperationalInsightSearchResult PowerShell cmdlet. Answer: A Explanation: The following query returns the transaction log size on each distribution. If one of the log files is reaching 160 GB, you should consider scaling up your instance or limiting your transaction size.-- Transaction log size SELECT instance_name as distribution_db, cntr_value*1.0/1048576 as log_file_size_used_GB, pdw_node_id FROM sys.dm_pdw_nodes_os_performance_counters WHERE instance_name like 'Distribution_%' AND counter_name = 'Log File(s) Used Size (KB)' References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-monitor> Question: 50 HOTSPOT You need to collect application metrics, streaming query events, and application log messages for an Azure Databricks cluster. Which type of library and workspace should you implement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Library:

Azure Databricks Monitoring Library
Microsoft Azure Management Monitoring Library
PyTorch
TensorFlow

Workspace:

Azure Databricks
Azure Log Analytics
Azure Machine Learning

Answer:

Library:

Azure Databricks Monitoring Library
Microsoft Azure Management Monitoring Library
PyTorch
TensorFlow

Workspace:

Azure Databricks
Azure Log Analytics
Azure Machine Learning

Explanation: You can send application logs and metrics from Azure Databricks to a Log Analytics workspace. It uses the Azure Databrick Monitoring Library, which is available on GitHub. References: <https://docs.microsoft.com/en-us/azure/architecture/databricks-monitoring/application-logs> Question: 51 You have a SQL pool in Azure Synapse. A user reports that queries against the pool take longer than expected to complete. You need to add monitoring to the underlying storage to help diagnose the issue. Which two metrics should you monitor? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point. A. Cache used percentage B. DWU Limit C. Snapshot Storage Size D. Active queries E. Cache hit percentage Answer: AE Explanation: A: Cache used is the sum of all bytes in the local SSD cache across all nodes and cache capacity is the sum of the storage capacity of the local SSD cache across all nodes. E: Cache hits is the sum of all columnstore segments hits in the local SSD cache and cache miss is the columnstore segments misses in the local SSD cache summed across all nodes Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-concept-resource-utilization-query-activity> Question: 52 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB. You plan to copy the data from the storage account to an Azure SQL data warehouse. You need to prepare the files to ensure that the data copies quickly. Solution: You modify the files to ensure that each row is more than 1 MB. Does this meet the goal? A. Yes B. No Answer: B Explanation: Instead modify the files to ensure that each row is less than 1 MB. References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data> Question: 53 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB. You plan to copy the data from the storage account to an Azure SQL data warehouse. You need to prepare the files to ensure that the data copies quickly. Solution: You modify the files to ensure that each row is less than 1 MB. Does this meet the goal? A. Yes B. No Answer: A Explanation: When exporting data into an ORC File Format, you might get Java out-of-memory errors when there are large text columns. To work around this limitation, export only a subset of the columns. References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data> Question: 54

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB. You plan to copy the data from the storage account to an enterprise data warehouse in Azure Synapse Analytics. You need to prepare the files to ensure that the data copies quickly. Solution: You convert the files to compressed delimited text files. Does this meet the goal? A. Yes B. No Answer: A Explanation: All file formats have different performance characteristics. For the fastest load, use compressed delimited text files. Reference: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data> Question: 55

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1. You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1. Solution: You use a dedicated SQL pool to create an external table that has an additional DateTime column. Does this meet the goal? A. Yes B. No Answer: A Question: 56 You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1. You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1. Solution: In an Azure Synapse Analytics pipeline, you use a data flow that contains a Derived Column transformation. A. Yes B. No Answer: B Question: 57 You have an Azure Synapse

Analytics dedicated SQL pool that contains a table named Table1. You have files that are ingested and loaded into an Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and Azure Data Lake Storage Gen2 container named container1. You plan to insert data from the files into Table1 and transform the data. Each row of data in the files will produce one row in the serving layer of Table1. You need to ensure that when the source data files are loaded to container1, the DateTime is stored as an additional column in Table1. Solution: In an Azure Synapse Analytics pipeline, you use a Get Metadata activity that retrieves the DateTime of the files. Does this meet the goal? A. Yes B. No Answer: B Question: 58 What should you recommend to prevent users outside the Litware on-premises network from accessing the analytical data store? A. a server-level virtual network rule B. a database-level virtual network rule C. a database-level firewall IP rule D. a server-level firewall IP rule Answer: A Explanation: Virtual network rules are one firewall security feature that controls whether the database server for your single databases and elastic pool in Azure SQL Database or for your databases in SQL Data Warehouse accepts communications that are sent from particular subnets in virtual networks. Server-level, not database-level: Each virtual network rule applies to your whole Azure SQL Database server, not just to one particular database on the server. In other words, virtual network rule applies at the server level, not at the database level. References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-vnet-service-endpoint-rule-overview> Question: 59 What should you do to improve high availability of the real-time data processing solution? A. Deploy identical Azure Stream Analytics jobs to paired regions in Azure. B. Deploy a High Concurrency Databricks cluster. C. Deploy an Azure Stream Analytics job and use an Azure Automation runbook to check the status of the job and to start the job if it stops. D. Set Data Lake Storage to use geo-redundant storage (GRS). Answer: A Explanation: Guarantee Stream Analytics job reliability during service updates Part of being a fully managed service is the capability to introduce new service functionality and improvements at a rapid pace. As a result, Stream Analytics can have a service update deploy on a weekly (or more frequent) basis. No matter how much testing is done there is still a risk that an existing, running job may break due to the introduction of a bug. If you are running mission critical jobs, these risks need to be avoided. You can reduce this risk by following Azure's paired region model. Scenario: The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability> Question: 60 What should you recommend using to secure sensitive customer contact information? A. data labels B. column-level security C. row-level security D. Transparent Data Encryption (TDE) Answer: B Explanation: Scenario: All cloud data must be encrypted at rest and in transit. Always Encrypted is a feature designed to protect sensitive data stored in specific database columns from access (for example, credit card numbers, national identification numbers, or data on a need to know basis). This includes database administrators or other privileged users who are authorized to access the database to perform management tasks, but have no business need to access the particular data in the encrypted columns. The data is always encrypted, which means the encrypted data is decrypted only for processing by client applications with access to the encryption key. References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-security-overview> Question: 61 HOTSPOT Which Azure Data Factory components should you recommend using together to import the daily inventory data from the SQL server to Azure Data Lake Storage? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



Answer: 1. Self-hosted integration runtime 2. Schedule trigger 3. Copy Activity Resources From: 1. 2021 Latest Braindump2go DP-203 Exam Dumps (PDF & VCE) Free Share: <https://www.braindump2go.com/dp-203.html> 2. 2021 Latest Braindump2go DP-203 PDF and DP-203 VCE Dumps Free Share:

https://drive.google.com/drive/folders/1iYr0c-2LfLu8iev_F1XZJhK_LKXNTGhn?usp=sharing 3. 2021 Free Braindump2go DP-203

Exam Questions Download: [https://www.braindump2go.com/free-online-pdf/DP-203-PDF\(34-44\).pdf](https://www.braindump2go.com/free-online-pdf/DP-203-PDF(34-44).pdf)

[https://www.braindump2go.com/free-online-pdf/DP-203-PDF\(34-44\).pdf](https://www.braindump2go.com/free-online-pdf/DP-203-PDF(34-44).pdf)

[https://www.braindump2go.com/free-online-pdf/DP-203-PDF-Dumps\(1-16\).pdf](https://www.braindump2go.com/free-online-pdf/DP-203-PDF-Dumps(1-16).pdf)

[https://www.braindump2go.com/free-online-pdf/DP-203-VCE\(45-55\).pdf](https://www.braindump2go.com/free-online-pdf/DP-203-VCE(45-55).pdf)

[https://www.braindump2go.com/free-online-pdf/DP-203-VCE-Dumps\(17-33\).pdf](https://www.braindump2go.com/free-online-pdf/DP-203-VCE-Dumps(17-33).pdf)Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!